

You can't outsource understanding.

Assessment should reveal the knowledge, reasoning, judgment, and skill students have developed. AI-supported work may contribute to the evidence, but it may not by itself establish learning when the outcome requires personal understanding or performance.

This guide is both a decision tool and an idea bank. It offers a framework for choosing evidence, then a set of named patterns and concrete variations faculty can adapt across disciplines, course types, modalities, and class sizes. The examples are meant to spark experimentation, not prescribe a single model. We will learn by trying, revising, and sharing what works.

Start with the learning outcome

Before choosing a format or an AI rule, identify what students must personally know, do, interpret, explain, or defend. Then decide what evidence would make that learning visible. A polished artifact may be useful evidence, but it is rarely sufficient by itself.

- Step 1.** Define the learning outcome in observable terms (and whether the demonstration is individual or team-based).
- Step 2.** Set the AI rule for each assignment stage: AI required, AI permitted, AI limited, or AI prohibited.
- Step 3.** Choose evidence that is traceable, defensible, or answerable, and set the conditions: open, secure, or staged.
- Step 4.** Pair patterns deliberately. The strongest designs usually draw on more than one kind of evidence.
- Step 5.** Match the evidence to the setting, stakes, and scale. Check modality, accessibility, confidentiality, licensing, and the information students may submit to AI systems.

Step 1. Define the learning outcome with Kelley Course Koach

Kelley Course Koach can pull an assignment directly from Canvas and apply Kelley's assessment framework to help identify what still works and what should change. Start at ai.kelley.iu.edu/Chat. Write the outcome so someone could watch a student demonstrate it, and name whether the demonstration is individual or team-based. **Vague:** understand discounted cash flow valuation. **Observable:** build a DCF model for an assigned firm, justify the discount rate and growth assumptions, and adapt the valuation when one assumption changes.

Use Kelley Course Koach suggestions as a design draft. Faculty remain responsible for the learning outcome, AI rules, accessibility, confidentiality, rubric, and final assessment decisions. The verbs carry the assessment: know, build, interpret, explain, defend, adapt. A team deliverable rarely establishes individual mastery on its own.

Step 2. Set AI rules for each assignment stage

Students need the AI rule stated for each major stage of an assignment, not just for the course as a whole, with every restriction tied to a reason. Mark each stage with one of four levels. The descriptions below are written to students, so you can paste them directly into a syllabus or assignment sheet:

AI required	Working with AI is part of what the assignment teaches. Use it for the stated stages or functions.
AI permitted	You choose whether and how to use it within the stated boundaries. Your reasoning and judgment are what is graded.
AI limited	Only specified functions are allowed, such as brainstorming but not drafting, or code explanation but not code generation.
AI prohibited	The assignment stage builds or verifies something you must hold personally. The restriction protects your learning.

One caution: AI limited is the hardest rule to verify and the most likely to generate disputes. Use it only when the boundary is clear and observable. Do not treat prompt logs as proof of compliance. Pair a limited rule with traceable evidence or convert it to AI permitted with required disclosure of how AI was used.

Step 3. Choose the evidence and the conditions

Element	Core question	Useful evidence
Traceable	Can the instructor see how the work developed and where judgment entered the process?	The game tape: checkpoints, revision memos, annotated drafts, workflow maps, version history, testing notes, or brief process reflections.
Defensible	Can the student justify choices, evidence, assumptions, and tradeoffs in context?	Decision rationales, source verification, comparison of alternatives, constraints, audience analysis, or critique and revision of AI output.
Answerable	Can the student explain, adapt, or apply the work when questioned or when conditions change?	Oral defense, live application, targeted follow-up, model walkthrough, new constraint, stakeholder challenge, or timed addendum.

These labels describe evidence type, not quality; alignment, accessibility, scoring consistency, and proportionality still matter. Alongside the evidence, set the conditions under which each stage is produced. **Open AI-enabled work** permits or requires specified AI use and evaluates how students direct, verify, revise, and integrate it. **A secure individual demonstration** restricts AI and unauthorized assistance while preserving approved accommodations, so the student personally demonstrates the target capability. **A staged combination** uses open development followed by a secure or controlled demonstration.

No single approach is inherently superior. The learning outcome, stakes, modality, and professional context determine the right mix. When personal mastery is central and the stakes warrant it, include at least one secure individual demonstration. “Secure” or “controlled” describes the assessment conditions, not a required proctoring technology.

Pillar 1: Making the Process Visible (Traceable)

These approaches shift attention from the final deliverable to the work as it unfolds. They ask students to document reasoning, testing, revision, and collaboration.

1. Watching the Game Tape

Best when you need to see how the student developed or revised the work, not just where it ended.

Variation	AI rule and condition	Setting and stakes	Design idea
Prompt and checkpoint annotations	AI permitted or required; open work	Individual or team; asynchronous; low to mid stakes	Students submit selected prompts, queries, or interaction checkpoints, with brief annotations explaining what changed, why, and what improved as a result.
Tracked revisions	AI permitted or required; open work	Individual or team; asynchronous or hybrid; mid stakes	Students submit tracked changes, annotated revisions, or a brief revision memo showing where they accepted, rewrote, or rejected AI-assisted text and why.
Screen-share walkthrough	AI permitted during development; individual explanation required	Individual; asynchronous; high stakes	Students submit a final coding project, financial model, or other substantial deliverable with a brief screen-share video. They walk through the work, explain the hardest technical or strategic hurdle they encountered, and describe how they navigated it.

2. Managing the AI Assist

Best when responsible delegation, quality control, and retained human judgment are part of the learning outcome.

Variation	AI rule and condition	Setting and stakes	Design idea
Workflow and quality-control map	AI permitted or required; open work	Individual or team; asynchronous; mid to high stakes	Students map the workflow for a major assignment, explain which tasks they delegated to AI, how they quality-checked that work, and which responsibilities they retained for their own judgment.

Variation	AI rule and condition	Setting and stakes	Design idea
Delegation audit	AI permitted or required; open work	Individual or team; asynchronous; low to mid stakes	Students submit one AI output they relied on at a key stage, together with their corrections: the errors they caught, the checks they ran, and what they changed before using it. The rubric rewards the quality of the verification, not the polish of the output.

Do not grade prompt volume, polished process narratives, or exhaustive logs. Grade the quality of decisions, verification, revision, and accountability.

Pillar 2: Owning the Judgment (Defensible)

These approaches place students in specific contexts with real constraints. They require students to verify accuracy, synthesize information, and show where human judgment adds value.

3. Not Just Any Case

Best when the task should resist generic answers by grounding the decision in a real audience, organization, event, or constraint.

Variation	AI rule and condition	Setting and stakes	Design idea
Firsthand-observation strategy	AI permitted or limited; open work	Individual or team; any modality; mid to high stakes	Students pitch a strategy for a campus operation or public-facing service they can study directly (a dining hall, the campus bookstore, a bus route, a recreation facility), incorporating constraints they could only learn through firsthand observation: hours, traffic patterns, pricing, layout, staffing visible from the floor. Students may use AI only on non-sensitive, de-identified observation notes. They must not photograph, record, quote, or identify individuals or include private information.
Disruption and revision	AI permitted in initial work; secure or controlled follow-up	Individual or team; staged, with synchronous follow-up; mid to high stakes	Students prepare a recommendation for a specific audience based on a case with realistic constraints. After a draft or interim analysis, they receive a supplier failure, policy change, labor complaint, market shock, or similar disruption. They revise the recommendation and explain what changed, what did not, and why.
Current-event stress test	AI permitted; open work with source verification	Individual or team; hybrid or asynchronous; mid to high stakes	Students prepare a recommendation for a specific audience and then incorporate a student-chosen current event that materially affects the decision. They explain how the new information changes the recommendation, what parts of the earlier analysis no longer hold, and what they had to verify for themselves rather than accept at face value.

4. Show Your Judgment

Best when students need to critique AI output or justify why they used, limited, or rejected AI.

Variation	AI rule and condition	Setting and stakes	Design idea
Evaluate the AI plan	AI required for the initial draft; controlled follow-up	Individual; staged with synchronous follow-up; high stakes	Students use AI to generate an initial business plan, then evaluate it against a course-specific rubric or decision framework. They identify the three most significant weaknesses, support the critique with independent evidence, explain why those weaknesses matter in context, and revise one section to show what stronger human judgment would require. In the synchronous follow-up, they defend the critique and the revision.
Choice with rationale	AI permitted	Individual; synchronous or asynchronous; mid stakes	Students are given a task and told they may use AI. They briefly justify the allocation in both directions: where they used AI and why, and where working manually produced a better result for the audience, context, or decision at hand.

Authentic assessment does not mean unrestricted AI use. Professional competence includes knowing when to use tools, how to supervise them, when to verify, and when to work without them.

Pillar 3: Responding and Adapting (Answerable)

These approaches require students to explain, interpret, or adapt their thinking when questioned or when conditions change. They help reveal whether students have internalized the work or are relying too heavily on outsourced reasoning. Live formats can also measure speed, composure, and verbal fluency, whether intended or not. Grade those capabilities only when they are part of the learning outcome; otherwise, offer an equivalent format without lowering the mastery bar.

5. Talk Me Through It

Best when students need to explain and defend work they prepared, including in a professional conversation.

Variation	AI rule and condition	Setting and stakes	Design idea
Targeted oral defense	AI permitted in preparation; prohibited during defense	Individual; in-person or synchronous online; high stakes	After submitting a major project, students complete a brief follow-up in which the instructor asks targeted questions, such as: "If you had 10% less budget for this recommendation, which element would you cut and why?"
Resistant stakeholder	AI permitted in preparation; prohibited during live exchange	Individual or team; in-person or synchronous online; mid to high stakes	Students prepare a sales pitch, recommendation, or proposal in advance. A resistant client or stakeholder pushes back on price, assumptions, risks, or fit. Students respond in real time, adjust their pitch, and defend their reasoning without AI assistance in the room.
The classroom becomes the boardroom	AI permitted in preparation; prohibited during live defense	Individual or team; in-person or synchronous online; mid to high stakes	Students may use AI to help prepare a memo, recommendation, or project. In the classroom, they must verbally defend the recommendation, respond to questions, and justify their decisions in real time without AI assistance.
Leadership review	AI permitted in development; secure live review	Individual or team; in-person or synchronous online; mid to high stakes	Students build a dashboard, model, or other decision tool in advance. The classroom becomes a meeting with leadership: students answer supervisor- or executive-style questions about the tool's logic, assumptions, limitations, and implications for decision-making. A weak tool, or weak understanding of the tool, becomes visible quickly.
Randomized recorded follow-up	AI prohibited during response; controlled individual demonstration	Individual; asynchronous; mid stakes	Students receive one or two questions selected from a bank and record a concise response within a defined window. The questions ask them to explain or adapt their own submission, not recite generic content.

6. Now Change One Thing

Best when the learning outcome includes adaptation, transfer, or judgment under changing conditions.

Variation	AI rule and condition	Setting and stakes	Design idea
Competitor price shock	AI permitted in preparation; prohibited during live adaptation	Individual or team; in-person or synchronous online; mid to high stakes	Students prepare a strategic plan as homework. In class, the instructor introduces a breaking development, such as a major competitor cutting prices by 20 percent. Students then have 30 minutes to revise their strategy.
Timed asynchronous addendum	AI prohibited or limited during follow-up; controlled window	Individual; asynchronous; low to mid stakes	After the main submission, students receive one new constraint during a short scheduled assessment window and submit a short addendum explaining what changes, what stays the same, and why.
New stakeholder concern	AI permitted in preparation; prohibited during live adaptation	Team; in-person or synchronous online; high stakes	During a live team presentation, the instructor introduces a new but plausible stakeholder concern, and the team must adjust its pitch in real time to address that perspective. Add a proportionate individual check when individual mastery matters.
Fresh team problem	AI prohibited during the response; controlled team demonstration	Team; in-person or synchronous online; mid stakes	Teams receive a fresh problem in class and have 45 minutes to map a response on a whiteboard or shared board. The rubric focuses on prioritization, tradeoffs, participation, and reasoning. Add a short individual follow-up when individual mastery matters.

Step 4. Pair patterns deliberately

Any single piece of asynchronous evidence, including revision memos, prompt annotations, and reflections, can itself be produced with AI. Documentation alone cannot establish individual mastery; an answerable moment anchors the rest. Across a course or program, the strongest designs therefore draw on more than one kind of evidence. These combinations are starting points, not prescriptions.

Learning claim	Open or developmental stage	Secure or controlled follow-up
Develop and defend a recommendation	Not Just Any Case: open AI-enabled analysis grounded in a real audience and constraints.	Talk Me Through It: targeted oral defense or resistant stakeholder challenge.
Use AI responsibly while retaining individual mastery	Managing the AI Assist: workflow and quality-control map for the team project.	Now Change One Thing: short individual adaptation under a new constraint.
Evaluate AI instead of accepting it	Show Your Judgment: critique and revise an AI-generated plan.	Fresh problem or targeted defense completed without AI assistance.
Build and understand a technical artifact	Watching the Game Tape: technical walkthrough, testing notes, and revision evidence.	Talk Me Through It: leadership review of the model, dashboard, code, or decision tool.

Step 5. Match the evidence to the setting, stakes, and scale

Context	Useful options	Design caution
In-person	Live application, short defense, whiteboard reasoning, role-play, or timed adaptation.	Provide accessible alternatives when a format creates a disability barrier.
Synchronous online	Screen-share walkthrough, oral questions, collaborative board, live case response, or stakeholder challenge.	Design for bandwidth, time zones, and technology failure.
Asynchronous	Staged submissions, revision memos, workflow maps, screencasts, randomized follow-ups, or short controlled addenda.	Do not mistake extensive logging for valid evidence of learning.
Hybrid	Open AI-enabled development plus an in-person or synchronous demonstration of the target capability.	Keep rules consistent across stages and explain any change in permitted tools.
Team work	Team process evidence, role allocation, shared defense, peer evidence, and a proportionate individual check.	A strong team product does not automatically establish individual mastery.
High stakes	Combine multiple forms of evidence, including at least one secure individual demonstration when personal mastery is central.	Use clear rubrics and reasonable notice. Avoid surprise as the main source of difficulty.
Low stakes	Frequent checks, short explanations, quick adaptations, retrieval practice, or reflection on AI errors.	Keep the documentation burden lower than the learning value.
Large enrollment	Rotating checks, individualized question banks, group demonstrations plus short individual responses, or brief recorded explanations.	Make expectations transparent. Do not let logistics replace valid evidence.

Grade the capability, not the polish

A rubric for AI-era evidence rewards the capability the evidence is meant to show, not the polish of the artifact. Separate the dimensions that matter. A useful rubric may score the artifact, disciplinary reasoning, evidence verification, AI-related judgment, and the ability to explain or adapt. Do not let surface polish compensate for weak understanding.

Evidence	Strong evidence looks like	Weak evidence looks like
Traceable	Checkpoints show substantive change over time, and annotations explain why revisions happened and what improved.	Logs are voluminous but generic; entries describe activity without reasoning.
Defensible	Choices are justified against alternatives, constraints, and audience; sources and AI output are verified.	Assertions appear without rationale; AI output is accepted at face value.

Evidence	Strong evidence looks like	Weak evidence looks like
Answerable	The student accurately explains and adapts their own submission under new questions or constraints.	Responses are generic, do not reference the actual submission, or collapse when one condition changes.

Communicate the assessment design

For each stage, name the AI category, permitted functions, required documentation, confidentiality limits, and the reason for the rule. Explain how the evidence will be evaluated. Students should know whether the instructor is assessing the final artifact, the process, the explanation, the adaptation, or a combination.

Example explanation to students

"This assessment asks you to demonstrate the learning outcome through both the work you produce and the decisions you can explain. AI is [required, permitted, limited, or prohibited] for the specified stages. You remain responsible for verifying the work and must be able to explain, adapt, and defend it."

This guide designs for evidence rather than detection. Detection tools alone cannot establish that a student violated an AI rule. When you suspect misuse, follow the suspected-misuse process in the Kelley Faculty AI Playbook rather than relying on a detector score.

Access, confidentiality, and proportionality

- When AI is required or materially advantageous, provide an IU-provided, course-provided, or genuinely comparable no-cost option, plus an equivalent route when the tool is unavailable or inaccessible. Refer students to [current IU-approved tools](#).
- Do not require students to submit teammates' work, client or employer information, NDA-protected material, licensed cases or databases, unpublished research, research-participant information, personal information, or credentials unless IU has approved the exact tool and use for that data.
- Use only the process evidence needed to support the learning claim. Avoid unnecessary surveillance, exhaustive prompt logs, or documentation that rewards compliance more than understanding.
- Provide an accessible alternative when a live, recorded, timed, or tool-dependent format creates a barrier. Plan for outages and technology failure.

Preflight checklist

- ☐ The learning outcome identifies what students must personally know, do, explain, or defend.
- ☐ The assignment uses one of the four AI categories at each major stage.
- ☐ The evidence is traceable, defensible, answerable, or a purposeful combination.
- ☐ Open AI-enabled work and secure individual demonstration are clearly distinguished.
- ☐ The design works for the course modality, stakes, team structure, accessibility needs, and scale.
- ☐ Students can complete the work without violating confidentiality, licenses, contracts, privacy requirements, or IU rules.
- ☐ The documentation burden is proportionate to the learning value.
- ☐ The rubric evaluates the targeted capability, not polish alone.
- ☐ Students know why the AI rule exists and which evidence will be evaluated.
- ☐ When personal mastery is central and the stakes warrant it, the design includes a secure individual demonstration.

Get help applying this guide

Kelley Course Koach can pull an assignment directly from Canvas and apply Kelley's assessment framework to help identify what still works and what should change. Start at ai.kelley.iu.edu/Chat.

Use its suggestions as a design draft. Faculty remain responsible for the learning outcome, AI rules, accessibility, confidentiality, rubric, and final assessment decisions.

Start small and share what you learn.

Redesign one assignment, test it, ask students what the structure made visible or confusing, and share the result with colleagues. The curriculum will evolve as faculty learn what works across disciplines, modalities, and class sizes.